

A METHODOLOGY FOR RISK-BASED SAFETY ASSURANCE IN THE MAJOR HAZARDS INDUSTRIES

P. Naylor BSc MSc PhD CEng CPhys CSci
Specialist Inspector, Human Factors, Health and Safety Executive
A. Taleb-Bendiab BSc PhD
Professor of Computing, School of Computing and Mathematical Sciences,
Liverpool John Moores University

There is a climate of change in the regulation of Health and Safety in the UK Major Hazards Industries as reflected in the Hampton Report[5], published in March 2005. The key recommendation in relation to regulators, made by Hampton, is that comprehensive risk assessment should form the basis of all regulatory activity. This paper outlines a *safety assurance* approach, in which assessment of the safety case drives the follow-on *verification* of arrangements in operation, for installations in the major hazards sectors. The approach (designed to be applicable beyond the manufacturing process sectors, potentially to any situation where safety assurance is required) consists of two complementary parts: the broad *methodology* with associated framework, and – within this – the *core model* for the analysis of risk and evaluation of protection measures.

The main area of focus, for this paper, is the *methodology*. We begin by defining terms and concepts adapted to the approach of the paper, and then proceed to outline the main functional elements and workflows involved: from analysis of the domain, classification of installations, specification of roles, knowledge modelling, hazard and risk analysis and protection evaluation, through to the specification of verification as an output of the risk assessment process. It is shown that empirical evidence on risk performance can be retained through a knowledge model to iteratively improve the predictive modelling of risk.

Using a Safety Case assessment-centred approach, we seek to integrate the activities of Safety and Risk Management, Risk analysis and Protection evaluation, and Human/Organisational Factors in such a way that the Safety Case can be used not just to determine the acceptability of risks and associated protection systems, but to both prioritise and specify follow-on operational verification in a consistent, structured and systematic manner. With this in mind, the aim and purpose of this paper, therefore, is to stimulate discussion on risk-based safety assurance as a collaborative *Dutyholder-Regulator* activity centred around a dynamic safety case, as clarified below.

INTRODUCTION

The current safety regulatory regime for the UK major hazards industries defines the *Dutyholder* as the body (company) with legal responsibility for safe operation of an installation. The *Regulator* with the "challenge" function under law, is the Health and Safety Executive within the context of this paper. The Safety Case regime for major hazards regulation in the UK centres upon demonstration, via a Dutyholder's safety case, that

an installation presents risks that can be managed within specified criteria, as a condition for commencement or continuation of operation. The specific criterion (rooted in the Health and Safety at Work Act, 1974[7]) is that risks are controlled by means of risk-reduction measures that reduce risk *So Far As Is Reasonably Practicable* (SFAIRP). The Safety Case, which is compiled and owned by the Dutyholder, is subsequently *assessed* by the HSE, as the regulator, through an evaluation of this demonstration. The regulatory counterpart to *assessment* is the inspection or, more exactly, the *verification* of the operational management of risks and associated risk-reduction measures, predominantly on-site at the installation.

Whilst methods have sought to draw upon assessment findings to inform verification activity, the two have – until recently – been practiced as largely separate processes, with relatively loose coupling. This position is now changing. The key driver of change is the Hampton Report[5] whose principal recommendation (Rec.1) is that all regulatory activity be based on clear, comprehensive risk assessment. It recommends a combined consideration of historic safety performance and predictive risk. Therefore, since the major hazards safety regulatory regimes in the UK are all safety case based, it is argued in this paper that the quality of risk assessment provided should aim to be sufficient to enable risk-based verification (inspection) activity to be implemented directly from it. This alludes to integrated assessment and verification.

To this end, this paper introduces the Safety Assurance Concept as an integrated approach to safety regulation in the major hazards industrial sector. Much of the detailed *core mode*l is drawn from the *Generic Risk and Protection Integration Model* (GRAPIM) developed at Liverpool John Moore's University, although the GRAPIM approach is not exclusively targeted at the industrial sector. Similarly, there is much in common in the broader methodology with the Australia & New Zealand risk management framework[1] standard, which may add value to the approach for risk-based verification.

To summarise the intention of this paper from the outset: to describe an outline methodology for the integrated assessment and verification associated with major hazards industrial installations. That is to say, verification informed by the safety case, by the formulation of a safety assurance concept that is strongly rooted in a formal method of integrated risk and protection analysis, defined as the core model within the overall methodology, discussed below.

THE SAFETY ASSURANCE METHODOLOGY

The approach outlined in this paper is one of Safety Assurance, which is analogous to the ISO9000 quality assurance concept. The essence of this is the determination of adequacy of the submission by *assessment*, and conformance of construction or practice to that stated in the submission by follow-on *verification*. In order to clarify this, it is necessary to define some terms within the context of this paper.

- ***Submission:*** the written safety case, or safety report, that provides the argument for safe design and/or operation of an installation.

- ***Safety Assurance:*** the combined activity of *assessment* of the submission and the follow-on *verification* of arrangements in the setting of the installation's operation.
- ***Adequacy:*** determination of sufficiency of content and particulars in the submission against the requirements of the regulations that apply (e.g. Offshore[8], or COMAH[10]).
- ***Conformance:*** determination of adherence to arrangements or particulars claimed in the submission.
- ***Evaluation:*** (in relation to arrangements or protection systems) the process of determination of fitness for purpose.
- ***Compliance:*** (in relation to evaluation) the adherence of arrangements to applicable regulation, code of practice, or guidance that apply.
- ***Performance:*** (in relation to evaluation) the process of determining the arrangements' or protection systems' risk-reducing effect.
- ***Assessment:*** the examination of information contained within the submission with the aims of determining *adequacy* of arrangements (SFAIRP), and the formulation of an active verification specification based on information contained within the submission.
- ***Verification:*** the examination of Dutyholder's safety arrangements, in relation to the installation's risks, in-situ at the installation, or directly with the Dutyholder.
- ***Active Verification:*** determination of *conformance* with claims or information made by the Dutyholder in the submission.
- ***Passive Verification:*** determination of quality of attributes of arrangements against regulations, standards, guidance or practice, not necessarily with reference to the submission.

In order to explain the methodology, refer to Figure 1. The methodology assumes the adaptation of this framework from first principles to a new domain or industry. It has more in common with the Australia and New Zealand standard for Risk Management[1] than with the UK Safety Management System guidance[9]. As can be seen, there are three main blocks: *preliminary*, *assessment* and *verification*.

Considering the *preliminary* block, there are four constituent workflows that are shown sequentially following the other. In reality, such workflows are more parallel than is shown. The first workflow is a domain analysis whereby the broad classes of installations are identified as those groups with sufficient common features in terms of risks and protection systems so as to form a broad template for the safety assurance activity for such a class, and thereby inform the allocation of specialist resources involved on a class-by-class basis. For example, within the offshore domain, we might consider the main structural classes of installation as being (for example): steel jacket sub-structure; concrete gravity sub-structure; floating semi-submersible; floating mono-hull; jack-up driller etc. We can then consider the operational variants within each of these structural classes: e.g. drilling only, production. Finally, we consider process variants: oil, gas, condensate etc. The net result of this Domain Analysis is an extensive, but finite, set of combinations of structural, operational and process variants to subsequently enable a more focused and detailed risk and protection analysis.

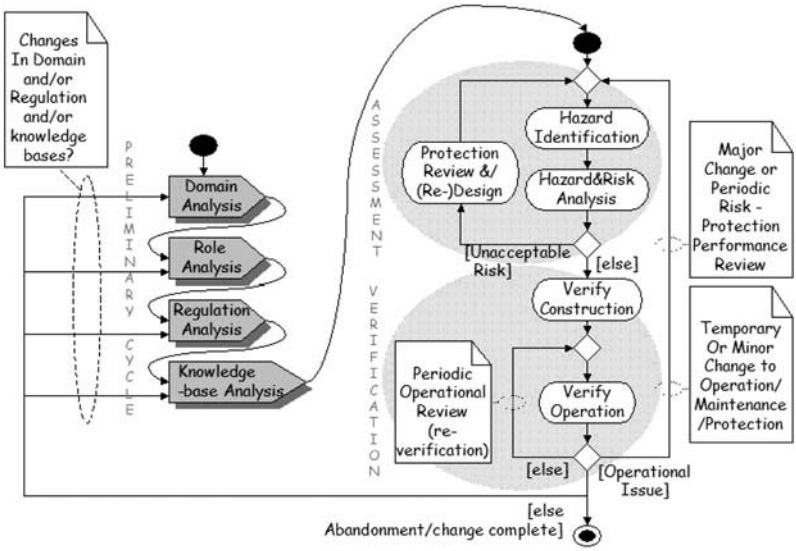

Figure 1. Methodology flowchart

Further in the paper, we will discuss the concept of *profiling*, as being the distribution and ranking of risk or protection significance in relation to an installation; however, the profile of specialisms (Figure 10) will be strongly associated with the installation class, as shown by the examples chosen: this is the aim of the role analysis workflow. Similarly, regulation analysis deals with the suite of applicable regulation sets that will vary depending upon the profile of disciplines and technology associated with a class, in a hierarchical manner (Figure 2). We refer to this as *Influence Base*, as the influence varies in priority from mandatory regulation, at one end, to preferred standards of practice in association with the protection technology, at the other. All of the above workflows may be combined, via knowledge analysis, to form a knowledge model that is incrementally developed through the process of assessment and verification which adds to the body of knowledge through what is termed *evolutionary growth*. The feedback loop for this is shown in the methodology flowchart of Figure 1, and specifically as an output of assessment, in Figure 3.

ASSESSMENT AND THE "CORE MODEL"

The Safety Assurance approach is essentially an assessment-centric one, in the sense that it is both enabled and facilitated by the preceding preliminary block of workflows described above, and then feeds forward into the subsequent verification activity block.

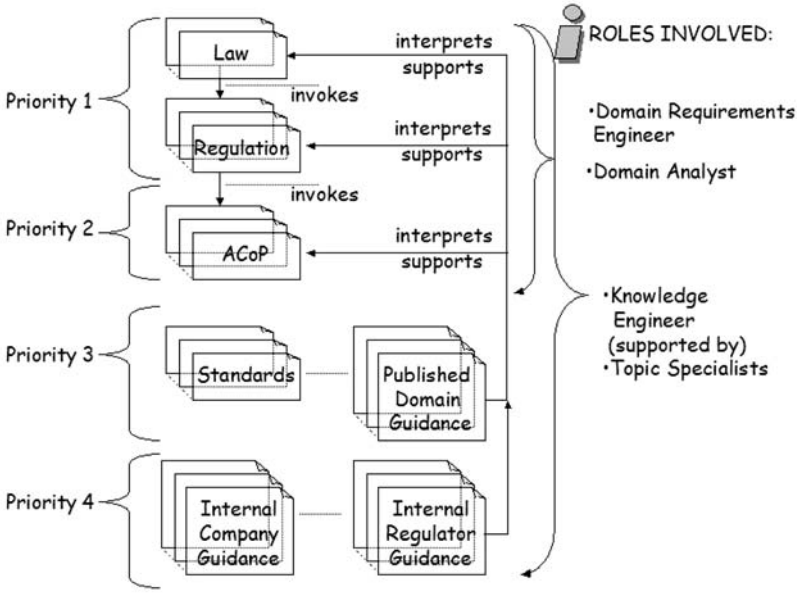

Figure 2. Influence base

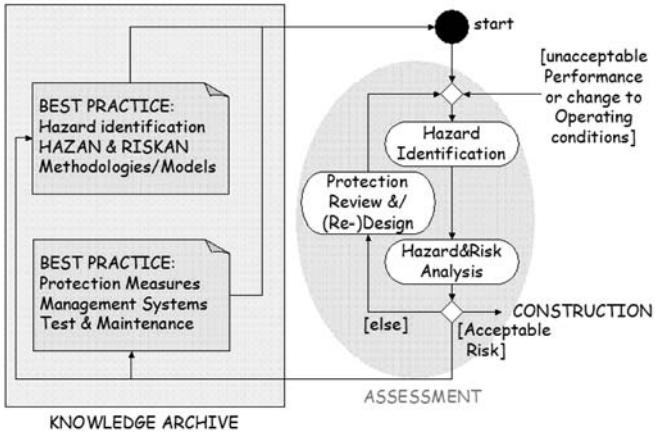

Figure 3. Simplified knowledge model

The enabling concept for the approach is a model of Risk Representation for an Installation, in which it is represented as the group of key hazard scenarios, modelled as risk "trees" centred on the *Critical Incidents* that define the hazards. Examples of such Critical Incidents (from an offshore setting) might be: loss of containment of hydrocarbon; dropped object over drill-deck; structural failure; ship collision etc. The key features of a critical incident are that it results from failure of systems designed to prevent the hazard from being realised (causes or *faults*), and that it presents a number of escalation paths via *events* that ultimately arrive at outcomes, whether hazardous or safe. A risk tree is shown in Figure 4. This is equivalent to the "bow-tie" model[6 and 17], but has a vertical root and branch orientation to assist with the form of data manipulation as described below.

If we look again at Figure 4, we can see that this form of cause-consequence model has a root structure that reflects the behaviour of protection measures or systems, the redundancy or backup features of which are embodied within the logic of the root system: failure of a protection system is, in effect, the *initiator* shown at the very base of the risk tree. Realisation of the *critical incident* depends upon the conditions of the root logic being satisfied. The series connection of protection systems is reflected in their coupling via "and" logic gates, and the parallel coupling by "or" logic gates. The *event* branches, by contrast, consist of the binary expansion from the *critical incident* via *event* gates toward outcomes with associated consequences. The event gates represent branching points that depend upon the successful effect (or otherwise) of control or mitigation measures. In other words, the fault-root system reflects the potential of failure of preventative measures associated with the hazard scenario, whereas the event branch

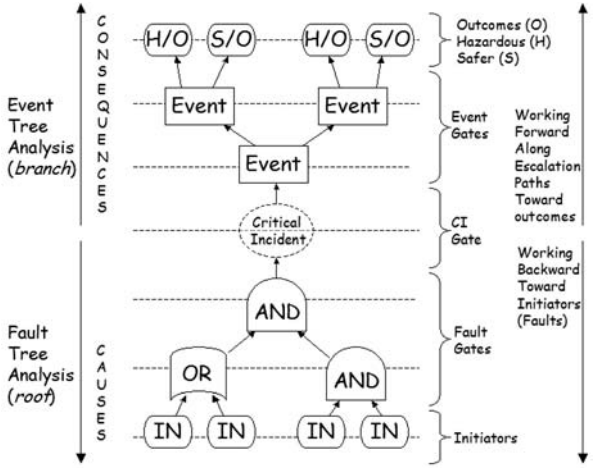

Figure 4. Fault-event tree

system represents the escalation paths associated with the failure of containment systems, defences or mitigation measures.

To complete the picture, the entire risk associated with an installation may be represented by a finite number of such risk-trees based on the critical incidents that are derived from the hazard identification activity within the risk assessment cycle. The numerical estimation of risk may be made by the summation of the risk associated with the outcomes of all such risk "trees".

This form of analysis is amenable to *Object-Oriented Analysis and Design (OOAD)*, the branch of computer science, or more specifically software engineering, whereby real-world entities may be identified and represented in terms of their attributes, operations, and interfaces with other objects within a working system. Although this technology has its origins in the commercial and telecommunications sectors, it is becoming an increasingly prevalent software modelling technique for the scientific and engineering communities, particularly with the development of the Unified Modelling Language (UML) as a standard for specification of Object-Oriented modelling.

The attributes and behaviours of risk trees and protection systems may be modelled effectively by this technique in such a way as to facilitate the calculation of risk and the impact of specific protection systems and technologies in relation to an installation. This is undertaken by generating unique identifiers (indices) for protection systems and their constituent protection duties (tasks) on the one hand, and by generating a similar set of unique indices for each gate within a risk tree, which itself has an index within the overall installation. This indexing arrangement is shown for a risk tree in Figure 5.

The model for safety assurance relies upon the creation of a persistent object model for a risk tree, where the type of object (the *class*) matches the type of *gate* within the risk tree. The class diagram showing the types of gate from which a risk-tree is constructed, is shown in Figure 7. Similarly, the different sub-types of protection are shown in the class diagram of Figure 8.

The inter-relation between protection-task objects and tree-gate objects within an installation is shown in Figure 6. The modelling process for an installation's risk and protection requires the analysis of critical incidents to represent them as a collection of gates, each of which is dynamically linked to a protection-task via a link-table that can be updated in the configuration process. At the most fundamental level, there is a one-to-one association between protection tasks and gates, via a dynamic *link-entry* within a *link-table* object. This allows flexible modelling of risk and protection performance in either design or risk analysis activities.

To clarify this point, it is important to consider a *protection-task* as influencing the output of a *gate* within a risk-tree. The key attribute of a protection-task is its probability of failure on demand (PFD) and its complement, both of which ascribe the probability of hazardous and safe outcomes of event gates within a tree. In the more detailed modelling of protection systems, reliability, availability, proof-test intervals etc., are attributes of the protection system and their constituent tasks, within which there can be internal (private in the UML) operations or methods within the protection object, in order to derive the PFD.

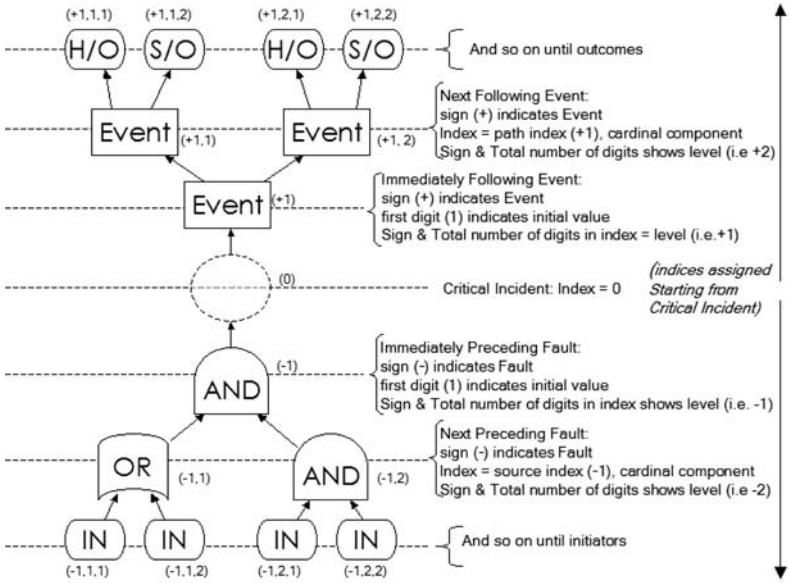

Figure 5. Tree with indices

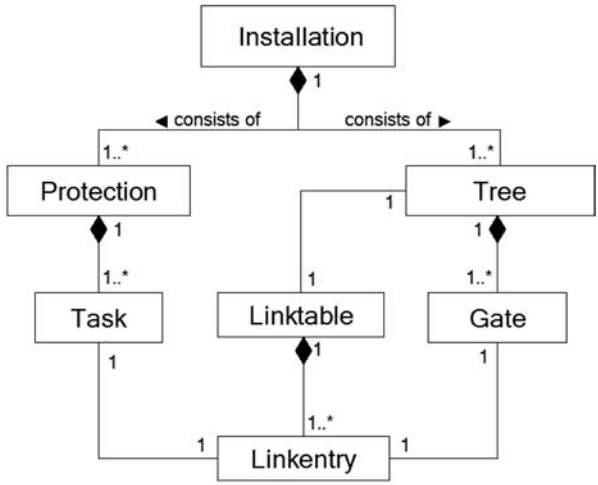

Figure 6. Tree object model

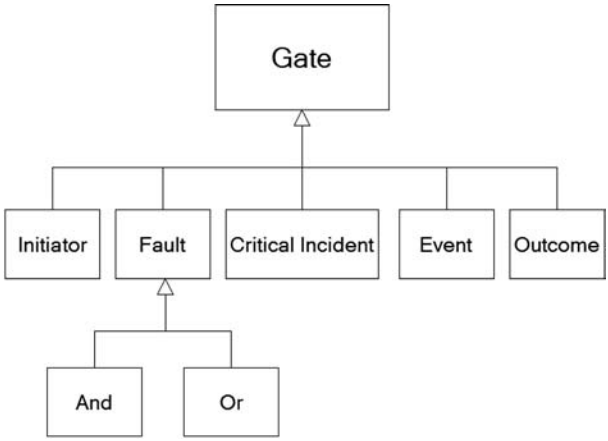

Figure 7. Gate class diagram

The operations (and therefore methods) that extract the PFD's from *protection* objects, are within the *risk-tree* objects, and the ultimate installation risk calculation operations are within the composite *installation* object. In brief, with this form of model, it is possible to rapidly derive risk-outcomes for an installation, repeatedly, on varying individual protection parameters, such as (e.g.) maintenance frequency, in order to assess their impact.

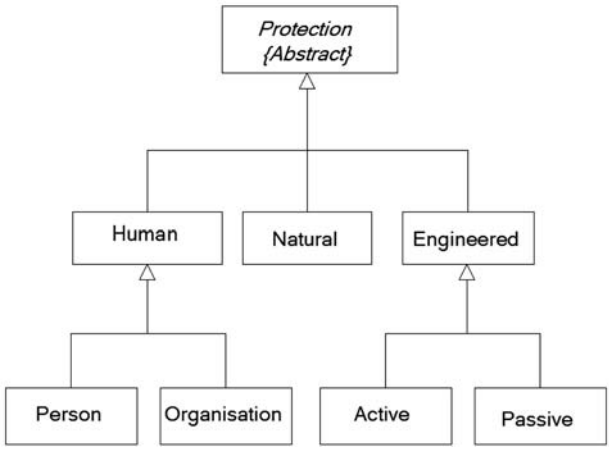

Figure 8. Protection class diagram

In addition, this particular object-oriented approach to risk analysis affords the possibility to ascribe a variety of attributes to both protection systems and risk-trees, in order to assist the Dutyholder and regulator alike. For example, we can attach *specialism*, *assessment*, or *verification* as attributes within a protection system function, in order to identify the regulatory specialist discipline for which it is of interest, and whether it is assessment or verification amenable, respectively.

OUTPUTS

The use of attributes and operations within the model, as described above, allows us to sort, evaluate, and rank risks and protection systems' impacts within an installation. By the use of attributes, we are able to sort by–for example–technology and specialist discipline required by the regulator to examine a particular hazard or protection system. Similarly, by ascribing the attribute of either *assessment* or *verification* to the protection system, the regulator is enabled to plan resource in relation to each of these activities.

To illustrate this point, consider Figures 9 and 10. Figure 9 depicts a hypothetical preliminary ranking of three particular classes of installation, by the individual risk presented in the safety case. This ranking system, which is consistent with one method of verification planning[2], presents a first step toward indicating the regulatory resource that may be required for a particular installation. However, the developed approach of this paper would enable this to be decomposed or refined in such a form as to indicate

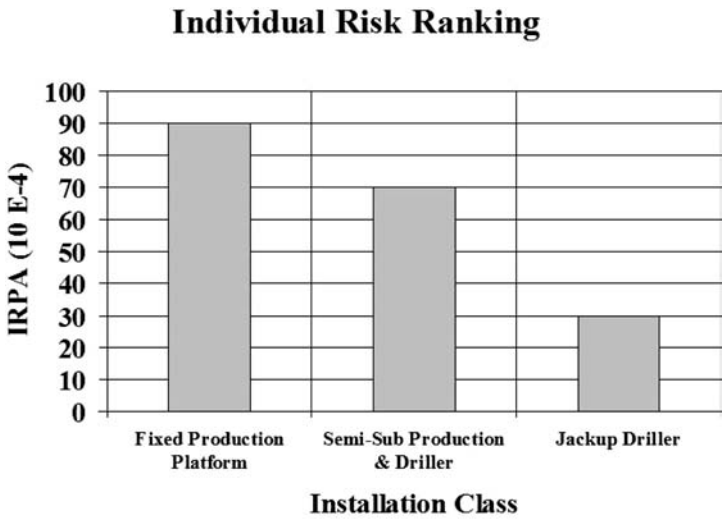

Figure 9. Illustration of aggregate risk profile

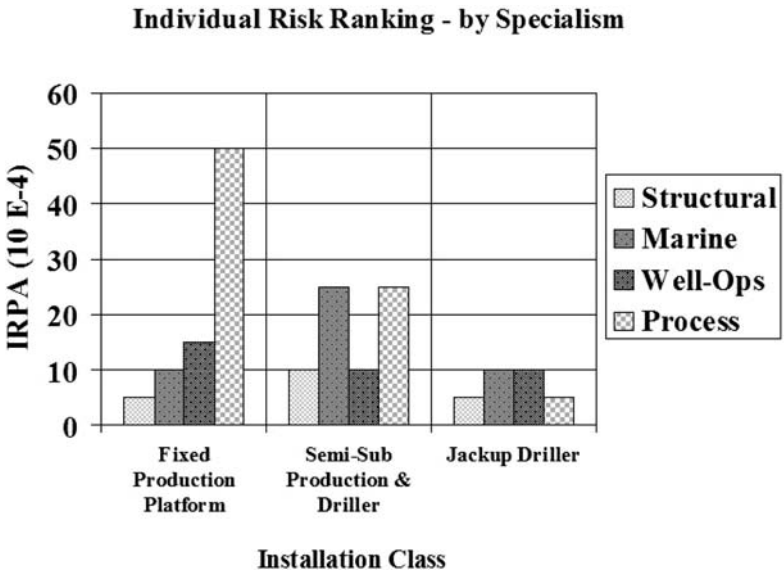

Figure 10. Illustration of decomposed risk profile

the constituent disciplines and, as a further development of this, in turn, the proportion of assessment and verification associated with the discipline (Figure 10).

This may, in large measure, satisfy the recommendation of Hampton[5] for a *"high quality risk assessment methodology and process* (that) *would enable the regulator to prioritise areas of work and use resources more effectively"* (para 4.39, p63).

VERIFICATION SPECIFICATION

Above and beyond the graphical profiling technique discussed above, the approach would enable the listing or scheduling of verification associated with an installation, by manipulation of those objects with verification attributes. It is, in addition, possible to rank such protection systems, according to their risk-reducing impact on the installation as a whole, and additionally, according to regulatory specialist discipline. This would add value to the regulatory process by providing clear structure and system to the organisation of regulatory resources.

FEEDBACK: COMPARISON OF OBSERVED WITH PREDICTED

So far, we have discussed the feedforward link from assessment to inform verification. However, there is an important feedback link from the verification process to our

knowledge base in relation to the installation to enable comparison of data gained through observation (verification) compared with that predicted in the safety case. Of specific interest, are frequencies of initiator-incidents, critical incidents or events; and also failure rates that may modify the reliability attributes of protection systems from those predicted.

It is possible, with the form of model described in this paper, to compare such values obtained by verification with those predicted in order to indicate discrepancy between predicted risk and that inferred by the verification process. This is shown on the feedback loop of Figure 1 to the Knowledge Analysis workflow.

How does this then inform the regulatory approach in relation to an installation or Dutyholder? In the first instance, it allows the data to be held for use in any subsequent revision of the safety case by the Dutyholder, in order to provide modified and more accurate prediction of risk. Secondly, where significant discrepancy appears between predicted and observed values, it may indicate greater regulatory attention.

This may–to some extent–address the need for a means to indicate where a regime of "earned autonomy" may be appropriate, although this aspect is obviously only one of a number of measurements that may be made and used to inform the level of autonomy a dutyholder may earn in relation to its installation.

CONSISTENCY WITH OTHER APPROACHES

The approach is a development of the broad risk ranking method[2] currently under consideration within the HSE offshore division. It takes the aggregated risk estimation for an installation, and decomposes it into constituent regulatory specialisms (termed *topics* within the HSE), in order to identify risk contribution by technology, and so appropriately target regulatory resource for assessment and verification activity.

Additionally, it is contended that the approach outlined in this paper is consistent with Layer of Protection Analysis[3] and may potentially facilitate use of this LOPA approach where un-mitigated *process* risk is compared with *mitigated* risk. This is performed by simply substituting PFD values of unity within risk trees, then reverting to the predicted PFDs as stored attributes provided by protection objects. In a similar manner the approach is consistent with the IEC61508 Safety Integrity Level determination concept[13]: at best it may assist the SIL determination process; at least, it does not conflict with this approach, as SIL is a mathematical function of PFD.

Furthermore, as the core model is a quantitative one, it is not considered to conflict with qualitative approaches such as the *Leading Performance Indicators*[16] approach, and may–in fact–complement this form of management system/culture measurement.

This sits well with maintenance determination methodologies, such as Reliability-Centred Maintenance[15]. This is because the comparison of observed (verification) aspects of protection systems with those predicted or planned is common to both this approach and that involved in RCM.

CONCLUSIONS AND RECOMMENDATIONS

The approach, in concept, seeks to address the requirement of Hampton[5] that all regulatory activity be on the basis of *"clear, comprehensive risk assessment"* (Rec. 1). It may, on further development, provide the recommended *"high quality risk assessment methodology and process* (that) *would enable the regulator to prioritise areas of work and use resources more effectively"* (para 4.39, p63), although further development would be necessary to achieve this.

However, one of the key differences in the approach proposed by this paper, is that the safety case itself should be (or become) the source of information and data that would in form the risk analysis process, and would become truly "dynamic not static" (Hampton, Rec. 1) as data gathered through verification is stored in a knowledge base to be incorporated into the safety case at periodic revision milestones.

There are, however, contentious issues: is the maintenance of a knowledge base best undertaken by Dutyholder, or regulator, or both, or by an alternative industry body? There is certainly a high level of collaboration required, and whilst Hampton (Rec. 3) alludes to a shift in emphasis from enforcement to advice, the question has to be asked whether this is feasible.

Nevertheless, this approach is proven in concept, so it is expected that the next stage of development is a larger scale case study, in order to test the practicality of the complete methodology in all its major sections: preliminary, assessment and verification, with an industrial scale risk-model, and prototype knowledge base and verification feedback cycle.

In summary, the approach seeks to integrate assessment and verification on the one hand, and risk and protection analysis on the other. It seeks to contextualise this process within a broader framework-methodology that classifies installations and establishes an influence and knowledge model. The net result is intended to satisfy the recommendations of the Hampton report's broad recommendations, whilst maintaining consistency with other more established methodologies. It is foreseeable that the goal of risk-based inspection is not an instantly achievable one; however, it is contended that this approach, alongside those also cited in this paper, may work to achieve the goal in an evolutionary manner.

The views expressed in this paper do not necessarily represent those of the Health and Safety Executive.

REFERENCES

1. Australia & New Zealand Standards (1999 and 2004): *AS/NZS 4360:2004 Risk Management.* Sydney, Australia: ISBN 0-7337-5904-1
2. Blackmore, A. and McIntosh, R. (2004). *Planning Specialist Inspection Work*. HSE Internal Paper.
3. Center for Chemical Process Safety – CCPS. (2001). *Layer of Protection Analysis – Simplified Process Risk Assessment.* American Institute of Chemical Engineers.
4. DoE. (1990). *The Public Enquiry into the Piper Alpha Disaster (the "Cullen Report"). P253-4.*London: Department of Energy / HMSO.

5. Hampton, P. (March 2005). *Reducing Administrative Burdens: Effective Inspection and Enforcement*. (Rec 1, p115; para 2.92, p43; para 4.39, p63). Norwich: HMSO.
6. Harte, H. (2005). *Bow Tie Diagrams with Traffic Lights as a Tool for Inspection Management*. HSE Internal Paper.
7. HSE (1974). *Health and Safety at Work etc. Act, 1974*. SI 1974/1439. The Stationary Office, 1974. ISBN 0 11 141439 X.
8. HSE (1992). *A Guide to the Offshore Installations (Safety Case) Regulations*. London: HMSO. Ref.L30.
9. HSE (1997). *Successful Health and Safety Management*. HS(G)65. Norwich: HMSO.
10. HSE (1999). A Guide to the Control of Major Accident Hazard Regulations 1999. Sheffield: HSE Books. Ref. L111.
11. HSE (2001a). *Reducing Risks, Protecting People.* Sheffield: HSE Books.
12. HSE (2005). *Guidance on the Assessment of Safety Cases.* (GASCET). HSE Internal Guidance Publication.
13. IEC (International Electrotechnical Commission) (1998 and 2002). *Standard IEC 61508: Functional safety of electrical/electronic/programmable electronic safety-related systems*. London: British Standards Institution.
14. International Standards Organisation – ISO (2000) BS-EN-9000-1. (1994 & 2000). *Model for Quality Assurance in Design, Development, Production, Installation and Servicing*. London: British Standards Institute.
15. Moubray, J. (1997). *Reliability-Centred Maintenance*. (2nd Edition) Oxford: Elsevier Butterworth-Heinemann.
16. Step Change in Safety. (2003). *Leading Performance Indicators: Guidance for Effective Use.* http://step.steel-sci.org.
17. Wilday, J. (2004). *Offshore Risk Assessment.* HSE Draft Internal Guidance Document.